



The Role of Personality and Linguistic Patterns in Discriminating Between Fake News Spreaders and Fact Checkers

Anastasia Giachanou¹(✉), Esteban A. Ríssola², Bilal Ghanem¹,
Fabio Crestani², and Paolo Rosso¹

¹ Universitat Politècnica de València, Valencia, Spain
angia9@upv.es, bigha@doctor.upv.es, proso@dsic.upv.es

² Università della Svizzera italiana, Lugano, Switzerland
{esteban.andres.rissola,fabio.crestani}@usi.ch

Abstract. Users play a critical role in the creation and propagation of fake news online by consuming and sharing articles with inaccurate information either intentionally or unintentionally. Fake news are written in a way to confuse readers and therefore understanding which articles contain fabricated information is very challenging for non-experts. Given the difficulty of the task, several fact checking websites have been developed to raise awareness about which articles contain fabricated information. As a result of those platforms, several users are interested to share posts that cite evidence with the aim to refute fake news and warn other users. These users are known as *fact checkers*. However, there are users who tend to share false information, who can be characterised as potential *fake news spreaders*. In this paper, we propose the CheckerOrSpreader model that can classify a user as a potential fact checker or a potential fake news spreader. Our model is based on a Convolutional Neural Network (CNN) and combines word embeddings with features that represent users' personality traits and linguistic patterns used in their tweets. Experimental results show that leveraging linguistic patterns and personality traits can improve the performance in differentiating between checkers and spreaders.

Keywords: Fact checkers detection · Personality traits · Linguistic patterns

1 Introduction

Although fake news, rumours and conspiracy theories exist for a long time, the unprecedented growth of social media has created a prosper environment for their propagation. Fake news are propagated rapidly in social media and indeed faster than real news [29]. Inaccurate and fabricated information can negatively influence users' opinions on different aspects, ranging from which political party to vote to doubting about the safety of vaccination. For example, research has

shown how medical misinformation can result to false treatment advice [17], whereas in the political domain, several researchers have underlined the influence of fake news on elections and referendums [2, 5].

Users play a critical role in all the different phases of the fake news cycle, from their creation to their propagation. However, users are dealing with an incredible huge amount of information everyday coming from different sources. Therefore, parsing this information and understanding if it is correct and accurate is almost impossible for the users who are non-experts. On the other side, experts such as journalists have the appropriate background to find relevant information and judge the credibility of the different articles and sources. In an attempt to raise awareness and inform users about pieces of news that contain fake information, several platforms (e.g., snopes¹, politifact², leadstories³) have been developed. These platforms employ journalists or other domain experts who thoroughly examine the claims and the information presented in the articles before they label them based on their credibility.

The advent of the fact checking platforms have resulted in a new type of social media users who have showed interest in halting the propagation of fake news. Users who consume and share news from social media can be roughly classified in the following two categories; (i) users that tend to believe some of the fake news and who further share them intentionally or unintentionally, characterised as *potential fake news spreaders*, and (ii) users who want to raise awareness and tend to share posts informing that these articles are fake, characterised as *potential fact checkers*⁴.

Even the detection of fake news has received a lot of research attention, the role of the users is still under-explored. The differentiation between checkers and spreaders is an important task and can further help in the detection of fake news⁵. This information can be further used by responsible recommendation systems to suggest to users that tend to share fake news, news articles from reliable sources in order to raise their awareness. Also, these systems should be regularly updated regarding the information they have for the users given that users can learn to better identify fake news.

We believe that checkers are likely to have a set of different characteristics compared to spreaders. For example, it is possible that checkers use different linguistic patterns when they share posts and have different personality traits compared to spreaders. We use the posts (i.e., tweets) of the users to extract a range of linguistic patterns and to infer their personality traits. We use Linguistic Inquiry and Word Count (LIWC) [18] to extract psychometric and linguistic style patterns of the posts and a vectorial semantics approach proposed by Neuman and Cohen [16] to infer the personality trait of the users.

¹ <https://www.snopes.com>.

² <https://www.politifact.com/>.

³ <https://leadstories.com/>.

⁴ For brevity we will refer to the users that have the tendency to share fake news as *spreaders* and to those that check the factuality of articles as *checkers*.

⁵ Here we should note that in this paper we focus only in the classification at a user level and we leave the exploration of the role of users at a post level as a future work.

The contributions of this paper can be summarised as follows:

- We create a collection that contains sets of tweets that are published by two different groups of social media users; users that tend to share fact check tweets (checkers) and those that tend to share fake news (spreaders).
- We extract different linguistic patterns and infer personality traits from the tweets posted by users to study their impact on classifying a user as a checker or spreader.
- We propose CheckerOrSpreader, a model based on a CNN network and hand-crafted features that refer to the linguistic patterns and personality traits, and which aims to classify a user as a potential checker or spreader.

The rest of the paper is organised as follows. Section 2 discusses related work on fake news detection. In Sect. 3 we present the collection and the process we followed to create it. Next we present the CheckerOrSpreader model in Sect. 4. Section 5 presents the evaluation process and the evaluation performance of the approach. Finally, Sect. 6 discusses the limitations and the ethical concerns regarding our study followed by the conclusions and future work in Sect. 7.

2 Related Work

The detection of fake news has attracted a lot of research attention. Among other problems, researchers have tried to address bot detection [22], rumour detection [21] and fact checking [7]. Many of the proposed works have explored a wide range of linguistic patterns to detect fake news such as the number of pronouns, swear words or punctuation marks. Rashkin et al. [23] compared the language of real news with that of satire, hoaxes, and propaganda based on features they extracted with the LIWC software [18]. Emotions and sentiment have been shown to play an important role in various classification tasks [6, 9]. In case of fake news, Vosoughi et al. [29] showed that they trigger different emotions than real news. In addition, Ghanem et al. [8] explored the impact of emotions in the detection of the different types of fake news, whereas Giachanou et al. [10] analysed the effect of emotions in credibility detection.

Users are involved in various steps in the life cycle of fake news, from creating or changing information to sharing them online. The tendency of some users to believe fake news depends on a range of different factors, such as network properties, analytical thinking or cognitive skills [20]. For example, Shu et al. [26] analysed different features, such as registration time and found that users that share fake news have more recent accounts than users who share real news. Vo and Lee [28] analysed the linguistic characteristics (e.g., use of tenses, number of pronouns) of fact checking tweets and proposed a deep learning framework to generate responses with fact checking intention.

The personality of the users is also likely to have an impact on the tendency of some users to believe fake news. A traditional way to measure the personality traits is via explicit questionnaires that persons are asked to fill. A number of researchers have employed those questionnaires and tried to find the relation

between personality traits and the use of social media [3,25] or information seeking behavior [12].

With all the advancements in Natural Language Processing, several studies have claimed that personality traits can also be inferred from the text generated by the user. In particular, several studies have addressed the problem of personality detection as a classification or a regression task based on text and conversations generated by the users [1,24]. In the present work, we use the posts that are written by users to extract linguistic patterns based on LIWC [18] and to infer their personality traits based a vectorial semantics approach proposed by Neuman and Cohen [16]. Differently from previous works, we explore the impact of those characteristics on classifying a user as a potential fake news spreader and fact checker based on the posts that he/she published.

3 Collection

There are different collections built in the field of fake news [27,28,30]. However, the majority of the previous datasets focus on the classification of the article as fake or not [27,30]. Vo and Lee [28] focus on fact checking but they collect fact check tweets and not previous tweets posted by the users. To the best of our knowledge, there is no collection that we can use for the task of differentiating users as checkers and spreaders. Therefore, we decided to build our own collection⁶. To build the collection, we first collect articles that have been debunked as fake from the Lead Stories website⁷. Crawling articles from fact check websites is the most popular way to collect articles since they are already labeled by experts. This approach has been already used by other researchers in order to create collections [27]. In total, we collected 915 titles of articles that have been labeled as fake by experts. Then, we removed stopwords from the headlines and we used the processed headlines to search for relevant tweets. Figure 1 shows the pipeline that we used to create the collection.

To extract the tweets we use Twitter API. In total we collected 18,670 tweets that refer to the articles from Lead Stories. For some of the articles we managed to collect a high number of tweets, whereas other articles were not discussed a lot in Twitter. Table 1 shows examples of the articles for which we collected the highest and lowest number of tweets. From this table, we observe that the most popular article was about a medical topic and for which we collected 1,448 tweets. In addition, Fig. 2 shows the number of collected tweets per article. We observe that the frequencies follow a heavy-tailed distribution since a lot of tweets were posted for few articles and very few tweets for a lot of articles.

The tweets that we collected can be classified in two categories. The first category contains tweets that debunk the original article by claiming its falseness (fact check tweet), and usually citing one of the fact checking websites (snopes, politifact or leadstories). The second category contains tweets that re-post the article (spreading tweet) implying its truthfulness. To categorise the tweets into

⁶ The collection and the code will be available upon acceptance.

⁷ <https://leadstories.com/>.

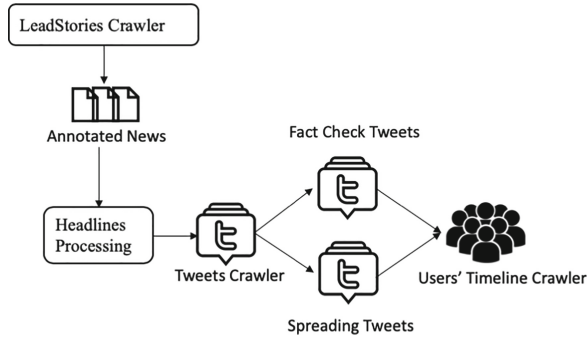


Fig. 1. Pipeline for the creation of the collection.

Table 1. Titles of the articles with the highest and lowest number of tweets.

Titles of articles with the highest number of tweets	Titles of articles with the lowest number of tweets
1. Doctors Who Discovered Cancer Enzymes In Vaccines NOT All Found Murdered	1. Make-A-Wish Did NOT Send Terminally Ill Spider-Man To Healthy Kid
2. Sugar Is NOT 8 Times More Addictive Than Cocaine	2. Man Did NOT Sue Radio Station For Playing Despacito 800 Times A Day
3. George H.W. Bush Did NOT Die at 93	3. Man-Eating Shark NOT Spotted In Ohio River
4. NO Alien Invasion This Is NOT Real	4. FBI DID NOT Classify President Obama As A Domestic Terrorist
5. First Bee Has NOT *Just* Been Added to Endangered Species List	5. Will Smith IS NOT Leaving America With His Family Never To Come Back

fact check and spreading tweets, we follow a semi-automated process. First, we manually identify specific patterns that are followed in the fact check tweets. According to those rules, if a tweet contains any of the terms {hoax, fake, false, fact check, snopes, politifact leadstories, lead stories} is a fact check tweet, otherwise it is a spreading tweet.

Figure 3 shows some examples of articles debunked as fake together with fact check and spreading tweets. We notice that in the fact check tweets we have terms such as fake, false and fact check, whereas in the spreading tweets we have re-posts of the specific article. Then, we manually checked a sample of the data to check if there are any wrong annotations. We manually checked 500 tweets and we did not find any cases of misclassification.

After the annotation of the tweets, we annotate the authors of the tweets as checkers or spreaders based on the number of fact check and spreading tweets they posted. In particular, if a user has both fact check and spreading tweets, then we consider that this user belongs to the category for which he/she has the larger number of tweets.




Fig. 2. Frequency distribution regarding of the number of tweets per article.

Finally, we collect the timeline tweets that the authors have posted to create our collection. In total, our collection contains tweets posted by 2,357 users, of which 454 are checkers and 1,903 spreaders.

4 CheckerOrSpreader

In this section, we present the CheckerOrSpreader system that aims to differentiate between checkers and spreaders. CheckerOrSpreader is based on a Convolutional Neural Network (CNN). The architecture of the CheckerOrSpreader system is depicted in Fig. 4.

CheckerOrSpreader consists of two different components, the word embeddings and the user’s psycho-linguistic component. The embeddings component is based on the tweets that users have posted on their timeline. The psycho-linguistic component represents the psychometric and linguistic style patterns and the personality traits that were derived from the textual content of the posts.

To extract the linguistic patterns and the personality traits we use the following approaches:

- **Linguistic patterns:** For the linguistic patterns, we employ LIWC [18] that is a software for mapping text to 73 psychologically-meaningful linguistic categories⁸. In particular, we extract pronouns (I, we, you, she/he, they), personal concerns (work, leisure, home, money, religion, death), time focus (past, present, future), cognitive processes (causation, discrepancy, tentative, certainty), informal language (swear, assent, nonfluencies, fillers), and affective processes (anxiety).
- **Personality scores:** The Five-Factor Model (FFM) [13], also called the Big Five, constitutes the most popular methodology used in automatic personality research [15]. In essence, it defines five basic *factors* or *dimensions* of personality. These factors are:

⁸ For a comprehensive list of LIWC categories see: <http://hdl.handle.net/2152/31333>.

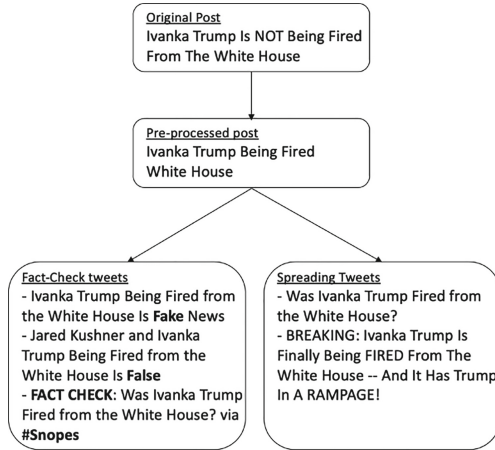


Fig. 3. Examples of fact check and spreading tweets.

- *openness to experience* (unconventional, insightful, imaginative)
- *conscientiousness* (organised, self-disciplined, ordered)
- *agreeableness* (cooperative, friendly, empathetic)
- *extraversion* (cheerful, sociable, assertive)
- *neuroticism* (anxious, sad, insecure)

Each of the five factors presents a *positive* and a complementary *negative* dimension. For instance, the complementary aspect to neuroticism is defined as *emotional stability*. Each individual can have a combination of these dimensions at a time. To obtain the personality scores, we followed the approach developed by Neuman and Cohen [16]. They proposed the construction of a set of vectors using a small group of adjectives, which according to theoretical and/or empirical knowledge, encode the essence of personality traits and personality disorders. Using a context-free word embedding they measured the semantic similarity between these vectors and the text written by different individuals. The similarity scores derived, allowed to quantify the degree in which a particular personality trait or disorder was evident in the text.

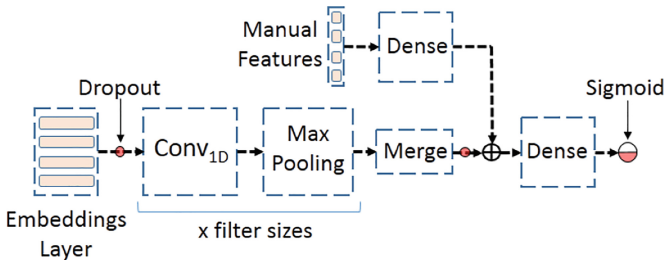


Fig. 4. Architecture of the CheckerOrSpreader model.

Table 2. Parameter optimisation for the different tested systems.

	Filters sizes	# of filters	Activation	Optimiser	Epochs
LSTM	64 (lstm units)		tanh	rmsprop	12
CNN	3,4	16	relu	adadelta	10
CNN+LIWC	3,4	32	relu	adadelta	15
CNN+personality	3,5	16	relu	adam	8
CheckerOrSpreader	4,5	32	relu	adam	13

5 Experiments

In this section we describe the experimental settings, the evaluation process and the results of our experiments.

5.1 Experimental Settings

For our experiments, we use 25% of our corpus of users for validation, 15% for test and the rest for the training. We initialize our embedding layer with the 300-dimensional pre-trained GloVe embeddings [19]. We allow the used embeddings to be tuned during the training process to fit more our training data. It's worth to mention that at the beginning of our experiments, we tested another version of our system by replacing the CNN with an Long Short-Term Memory (LSTM) network. The overall results showed that the CNN performs better for the particular task.

To find the best parameters of the different approaches on the validation set, we use the hyperopt library⁹. Table 2 shows the optimisation parameters for each approach.

5.2 Evaluation

For the evaluation, we use macro-F1 score. We use the following baselines to compare our results:

- *SVM+BoW* is based on Support Vector Machine (SVM) classifier trained on bag of words using Term Frequency - Inverse Document Frequency (Tf-Idf) weighting scheme.
- *Logistic Regression* trained on the different linguistic and personality scores features. In particular, we tried sentiment, emotion, LIWC and personality traits. For emotions we use NRC emotions lexicon [14] and we extracted anger, anticipation, disgust, fear, joy, sadness, surprise, and trust. We use the same lexicon to estimate the positive and negative sentiment in users' tweets.

⁹ <https://github.com/hyperopt/hyperopt>.

- *Universal Sentence Encoder (USE)* [4]: For the USE baseline, we represent the final concatenated documents (tweets) using USE embeddings¹⁰.
- *LSTM*: is based on a LSTM network with Glove pre-trained word embeddings for word representation.
- *CNN*: is a CNN with Glove pre-trained word embeddings for word representation.

5.3 Results

Table 3 shows the results of our experiments. We observe that CNN performs better than LSTM when they are trained only using word embeddings. In particular, CNN outperforms LSTM by 20.41%. Also, we observe that Logistic Regression achieves a low performance when it is trained with the different psycho-linguistic features. The best performance regarding Logistic Regression is achieved with the linguistic features extracted with LIWC.

Table 3. Performance of the different systems on the fact checkers detection task.

	F1-score
SVM+BoW	0.48
USE	0.53
LR+emotion	0.45
LR+sentiment	0.44
LR+LIWC	0.50
LR+personality	0.44
LSTM	0.44
CNN	0.54
CNN+LIWC	0.48
CNN+personality	0.57
CheckerOrSpreader	0.59

From Table 3 we also observe that combining CNN with the personality traits leads to a higher performance compared to combining CNN with the LIWC features. In particular, CNN+personality outperforms CNN+LIWC by 17.14%. This is an interesting observation that shows the importance of considering personality traits of users for their classification in checkers and spreaders.

Also, the results show that CheckerOrSpreader (CNN+personality+LIWC) achieves the best performance. In particular, CheckerOrSpreader manages to improve the performance by 8.85% compared to the CNN baseline and by 3.45% compared to the CNN+personality version.

¹⁰ <https://tfhub.dev/google/universal-sentence-encoder-large/3>.

6 Limitations and Ethical Concerns

Even if our study can provide valuable insights regarding the profile of spreaders and their automated detection, there are some limitations and ethical concerns. One limitation of our study is the use of an automated tool to infer the personality traits of the users based on the tweets that they have posted. Even if this tool has been shown to achieve good prediction performance, it is still prone to errors similar to all the automated tools. That means that some of the predictions regarding the personality traits that were inferred from the tweets might not be completely accurate. However, it is not possible to evaluate the performance of this tool on our collection since we do not have ground truth data regarding the users' personality traits. An alternative way to obtain information regarding the personality traits would be to contact these users and ask them to fill one of the standard questionnaires (e.g., IPIP questionnaire [11]) that have been evaluated based on several psychological studies and tend to have more precise results. However, the feasibility of this approach depends on the willingness of the users to fill the questionnaire.

Our study has also some ethical concerns. We should mention that the aim of a system that can differentiate between potential checkers and spreaders should be used by no means to stigmatise the users that have shared in the past fake news. On the contrary, such a tool should be used only for the benefit of the users. For example, it could be used as a supportive tool to prevent propagation of fake news and to raise awareness to users. We also want to highlight that a system that differentiates users to potential spreaders and checkers requires to consider ethics at all steps.

This study has also some ethical concerns regarding the collection and the release of the data. First, we plan to make this collection available only for research purposes. To protect the privacy of users, we plan to publish the data anonymized. Also, we plan to use neutral annotation labels regarding the two classes (i.e., 0 and 1 instead of checker and spreader) since we do not want to stigmatise specific users. Future researchers that want to use the collection will not have access to the information of which class each label refers to. Finally, we will not make available the labels at a post level since this information can reveal the information regarding the annotation labels at a user level.

7 Conclusions

In this paper, we focused on the problem of differentiating between users that tend to share fake news (spreaders) and those that tend to check the factuality of articles (checkers). To this end, we first collect articles that have been manually annotated from experts as fake or fact and then we detect the users on Twitter that have posts about the annotated articles. In addition, we propose the CheckerOrSpreader model that is based on a CNN network. CheckerOrSpreader incorporates the linguistics patterns and the personality traits of the users that are inferred from users' posts to decide if a user is a potential spreader or checker.

Experimental results showed that linguistic patterns and the inferred personality traits are very useful for the task.

In future, we plan to investigate how the linguistic and personality information that is extracted from users' posts can be incorporated into the systems that detect fake news.

Acknowledgements. The work of the first author is supported by the SNSF Early Postdoc Mobility grant under the project Early Fake News Detection on Social Media, Switzerland (P2TIP2_181441). The work of Paolo Rosso is partially funded by the Spanish MICINN under the research project MISMIS-FAKEHATE on Misinformation and Miscommunication in social media: FAKE news and HATE speech (PGC2018-096212-B-C31).

References

1. Bai, S., Zhu, T., Cheng, L.: Big-Five Personality Prediction Based on User Behaviors at Social Network Sites. <https://arxiv.org/abs/1204.4809> (2012)
2. Bastos, M.T., Mercea, D.: The Brexit botnet and user-generated hyperpartisan news. *Soc. Sci. Comput. Rev.* **37**(1), 38–54 (2019)
3. Burbach, L., Halbach, P., Zieffle, M., Calero Valdez, A.: Who shares fake news in online social networks? In: Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization, UMAP 2019, pp. 234–242 (2019)
4. Cer, D., et al.: Universal Sentence Encoder. <https://arxiv.org/abs/1803.11175> (2018)
5. DiFranzo, D., Gloria, M.J.K.: Filter Bubbles and Fake News. *ACM Crossroads* **23**(3), 32–35 (2017)
6. Farías, D.I.H., Patti, V., Rosso, P.: Irony detection in Twitter: the role of affective content. *ACM Trans. Internet Technol. (TOIT)* **16**(3), 1–24 (2016)
7. Ghanem, B., Glavaš, G., Giachanou, A., Paolo, S., Ponzetto, P.R., Rangel, F.: UPV-UMA at CheckThat! lab: verifying Arabic claims using a cross lingual approach. In: Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum (2019)
8. Ghanem, B., Rosso, P., Rangel, F.: An emotional analysis of false information in social media and news articles. *ACM Trans. Internet Technol. (TOIT)* **20**(2), 1–18 (2020)
9. Giachanou, A., Gonzalo, J., Crestani, F.: Propagating sentiment signals for estimating reputation polarity. *Inf. Process. Manage.* **56**(6), 102079 (2019)
10. Giachanou, A., Rosso, P., Crestani, F.: Leveraging emotional signals for credibility detection. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, pp. 877–880 (2019)
11. Goldberg, L.R.: A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. *Pers. Psychol. Europe* **7**(1), 7–28 (1999)
12. Heinström, J.: Five personality dimensions and their influence on information behaviour. *Inf. Res.* **9**(1), 1–9 (2003)
13. John, O.P., Srivastava, S.: The big-five trait taxonomy: history, measurement, and theoretical perspectives. In: Handbook of Personality: Theory and Research, pp. 102–138 (1999)

14. Mohammad, S.M., Turney, P.D.: Emotions evoked by common words and phrases: using mechanical turk to create an emotion lexicon. In: Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, pp. 26–34 (2010)
15. Neuman, Y.: Computational Personality Analysis: Introduction, Practical Applications and Novel Directions, 1st edn. Springer, Heidelberg (2016). <https://doi.org/10.1007/978-3-319-42460-6>
16. Neuman, Y., Cohen, Y.: A vectorial semantics approach to personality assessment. *Sci. Rep.* **4**(1), 1–6 (2014)
17. Oyeyemi, S.O., Gabarron, E., Wynn, R.: Ebola, Twitter, and misinformation: a dangerous combination? *BMJ Clin. Res.* **349**, g6178 (2014)
18. Pennebaker, J.W., Boyd, R.L., Jordan, K., Blackburn, K.: The Development and Psychometric Properties of LIWC 2015. Technical report (2015)
19. Pennington, J., Socher, R., Manning, C.: Glove: global vectors for word representation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, pp. 1532–1543 (2014)
20. Pennycook, G., Rand, D.: Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *J. Pers.* **88**, 185–200 (2018)
21. Qazvinian, V., Rosengren, E., Radev, D.R., Mei, Q.: Rumor has it: identifying misinformation in microblogs. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP 2011, pp. 1589–1599 (2011)
22. Rangel, F., Rosso, P.: Overview of the 7th author profiling task at PAN 2019: bots and gender profiling in Twitter. In: Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum (2019)
23. Rashkin, H., Choi, E., Jang, J.Y., Volkova, S., Choi, Y.: Truth of varying shades: analyzing language in fake news and political fact-checking. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 2931–2937 (2017)
24. Rissola, E.A., Bahrainian, S.A., Crestani, F.: Personality recognition in conversations using capsule neural networks. In: 2019 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2019, pp. 180–187 (2019)
25. Ross, C., Orr, E.S., Sisic, M., Arseneault, J.M., Simmering, M.G., Orr, R.R.: Personality and motivations associated with Facebook use. *Comput. Hum. Behav.* **25**(2), 578–586 (2009)
26. Shu, K., Wang, S., Liu, H.: Understanding user profiles on social media for fake news detection. In: Proceedings of the 2018 IEEE Conference on Multimedia Information Processing and Retrieval, MIPR 2018, pp. 430–435 (2018)
27. Shu, K., Mahudeswaran, D., Wang, S., Lee, D., Liu, H.: FakeNewsNet: A Data Repository with News Content, Social Context and Dynamic Information for Studying Fake News on Social Media. <https://arxiv.org/abs/1809.01286> (2018)
28. Vo, N., Lee, K.: Learning from fact-checkers: analysis and generation of fact-checking language. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, pp. 335–344 (2019)
29. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. *Science* **359**(6380), 1146–1151 (2018)
30. Wang, W.Y.: Liar, Liar Pants on Fire: A New Benchmark Dataset for Fake News Detection. <https://arxiv.org/abs/1705.00648> (2017)